

CLAIMS

We claim:

1. A computer system, comprising:
multiple processors,
5 a plurality of resources assigned to node groups;
a first descriptor of respective topological levels of at least one of the resources;
and
a second descriptor of respective performance of said resources.
- 10 2. The system of claim 1, wherein said first descriptor is a first level data structure,
and said second descriptor is a primary data structure.
3. The system of claim 1, wherein said primary data structure comprises a pointer to
a secondary data structure.
4. The system of claim 1, further comprising a node identifier for each node for
identifying positional placement of a resource.
- 15 5. The system of claim 4, wherein said node identifier represents multiple levels of
interconnect.
6. The system of claim 1, further comprising a dynamic updator of at least the first
and second descriptors.
- 20 7. The system of claim 6, wherein said dynamic updator reflects real-time system
configuration into the first descriptor.

8. The system of claim 6, wherein said dynamic updater reflects real-time system performance into the second descriptor.
9. The system of claim 1, wherein said second descriptor is selected from the group consisting of: processor descriptors, bus descriptors, memory descriptors and shared cache descriptors.
10. The system of claim 9, wherein said shared cache descriptor reflects interconnect of the system.
11. The system of claim 10, wherein said shared cache descriptor reflects latencies of the interconnects.
12. The system of claim 1, wherein said first descriptor reflects average latency between the node groups.
13. An article comprising:
a computer-readable signal bearing medium readable by a computer having multiple processors and
a plurality of resources assigned to node groups;
means in the medium for determining topological levels of at least some of the resources; and
means in the medium for determining performance of said resources.
14. The article of claim 13, wherein the medium is a recordable data storage medium.
15. The article of claim 13, wherein the medium is a modulated carrier signal.
16. The article of claim 13, wherein said topological level determining means is a first descriptor and said performance determining means is a second descriptor.

17. The article of claim 13, further comprising a node identifier for identifying positional placement of a resource for each node.
18. The article of claim 16, wherein said second descriptor is selected from the group consisting of: processor descriptors, bus descriptors, memory descriptors and shared cache descriptors.
19. The article of claim 13, wherein said second descriptor comprises a shared cache descriptor which reflects interconnect of resources.
20. The article of claim 19, wherein said shared cache descriptor reflects latencies of the interconnects.
21. The system of claim 16, wherein said second descriptor reflects average latencies between node groups.
22. A method for enabling allocation of resources in a multiprocessor, comprising: assigning multiple resources into node groups; and maintaining system resource topology and performance descriptions as at least one data structure.
23. The method of claim 22, further comprising traversing the data structure to enable allocation of at least some of the resources..
24. The method of claim 22, wherein said traversal step includes accessing a second data structure.
25. The method of claim 24, wherein said secondary data structure is selected from the group consisting of: processor descriptors, bus descriptors, memory descriptors and shared cache descriptors.

26. The method of claim 24, wherein said secondary data structure includes a shared cache descriptor for describing at least part of a system interconnect including latency between sibling nodes.

27. The method of claim 22, further comprising maintaining at least average latency between at least two of the nodes.

28. The method of claim 22, wherein said traversal step includes recursively accessing additional data structure levels.